

Claims

What is claimed is:

1. A method for executing database queries including identifying redundant conditions, comprising the steps of:
 - 5 (a) identifying a first set of conditions corresponding to a selected step for executing a query;
 - (b) identifying a second set of conditions corresponding to one or more steps for executing the query that feed the selected step;
 - (c) for each condition in the first set, checking whether the condition is mathematically
10 redundant, including redundant without being equivalent, with the other conditions in the union of the conditions corresponding to the selected step and the conditions in the second set,
 - (d) including each condition in the first set that is redundant as checked in step (c) in a
15 third set; and
 - (e) if there is only one condition in the third set after performing step (d), storing an identifier of the one condition.
2. The method of claim 1, further comprising the step of performing a cardinality calculation without considering conditions with stored identifiers.
- 20 3. The method of claim 1, further comprising the steps of:
 - (f) if there are a plurality of conditions in the third set after performing step (d), performing steps (g) and (h);
 - (g) sorting the conditions in the third set to identify an initial condition; and
 - (h) storing an identifier of the initial condition.
- 25 4. The method of claim 1, wherein step (e) includes ceasing storage of identifiers of additional conditions and the union in step (c) does not include conditions with stored identifiers, and further comprising the steps of:
 - (f) if there are no conditions in the third set after performing step (d), ceasing storage of identifiers of additional conditions;

- (g) if there are a plurality of conditions in the third set after performing step (d), performing steps (h) - (k);
- (h) sorting the conditions in the third set to identify an initial condition;
- (i) storing an identifier of the initial condition;
- 5 (j) modifying the first set of conditions to include only the conditions in the third set that are not the initial condition; and
- (k) starting again at step (c).

5. The method of claim 4, further comprising the step of performing a cardinality calculation without considering conditions with stored identifiers.

10 6. The method of claim 1 where the one or more steps for executing the query that feed the selected step include all the steps that feed the selected step.

7. The method of claim 1 further comprising the step of executing the query without applying conditions with stored identifiers.

8. The method of claim 1 further comprising the steps of:

- 15 (f) for each condition with a stored identifier that is a single table condition, performing steps (g)-(i);
- (g) if the condition is an equality with an index of the single table, identifying that condition as CPU effective;
- (h) if the condition is a range comparison of a value ordered index of the single table,
- 20 identifying that condition as CPU effective;
- (i) if the condition reduces the size of a temporary result in the execution of the query, identifying that condition as CPU effective; and
- (j) executing the query without applying conditions with stored identifiers that are not identified as CPU effective.

25 9. The method of claim 8 where identifying a condition as CPU effective comprises not identifying the condition as CPU redundant.

10. The method of claim 1 further comprising the steps of:

- (f) for each condition with a stored identifier that is a join condition, performing step (g);

- (g) if the condition is useful in co-locating the sources of the join, identifying that condition as CPU effective;
- (h) executing the query without applying conditions with stored identifiers that are not identified as CPU effective.

5 11. A computer program, stored on a tangible storage medium, for executing database queries, the program including executable instructions that cause a computer to:

- (a) identify a first set of conditions corresponding to a selected step for executing a query;
- 10 (b) identify a second set of conditions corresponding to one or more steps for executing the query that feed the selected step;
- (c) for each condition in the first set, check whether the condition is mathematically redundant, including redundant without being equivalent, with the other conditions in the union of the conditions corresponding to the selected step and
15 the conditions in the second set,
- (d) include each condition in the first set that is redundant as checked in step (c) in a third set; and
- (e) if there is only one condition in the third set after performing step (d), store an identifier of the one condition.

20 12. The program of claim 11, further including executable instructions that cause the computer to perform a cardinality calculation without considering conditions with stored identifiers.

13. The program of claim 11, further including executable instructions that cause the computer to:

- 25 (f) if there are a plurality of conditions in the third set after performing step (d), perform steps (g) and (h);
- (g) sort the conditions in the third set to identify an initial condition; and
- (h) store an identifier of the initial condition.

14. The program of claim 11, wherein step (e) includes ceasing to store identifiers of additional conditions and the union in step (c) does not include conditions with stored identifiers, and further including executable instructions that cause the computer to:

- (f) if there are no conditions in the third set after performing step (d), cease to store
5 identifiers of additional conditions;
- (g) if there are a plurality of conditions in the third set after performing step (d), perform steps (h) and (i);
- (h) sort the conditions in the third set to identify an initial condition;
- (i) store an identifier of the initial condition;
- 10 (j) modify the first set of conditions to include only the conditions in the third set that are not the initial condition; and
- (k) start again at step (c).

15. The program of claim 14, further including executable instructions that cause the computer to perform a cardinality calculation without considering conditions with stored identifiers.

15 16. The program of claim 11 where the one or more steps for executing the query that feed the selected step include all the steps that feed the selected step.

17. The program of claim 11 further including executable instructions that cause the computer to execute the query without applying conditions with stored identifiers.

18. The program of claim 11 further including executable instructions that cause the computer
20 to:

- (f) for each condition with a stored identifier that is a single table condition, perform steps (g)-(i);
- (g) if the condition is an equality with an index of the single table, identify that condition as CPU effective;
- 25 (h) if the condition is a range comparison of a value ordered index of the single table, identify that condition as CPU effective;
- (i) if the condition reduces the size of a temporary result in the execution of the query, identify that condition as CPU effective; and

(j) execute the query without applying conditions with stored identifiers that are not identified as CPU effective.

19. The program of claim 18 where identifying a condition as CPU effective comprises not identifying the condition as CPU redundant.

20. The program of claim 11 further including executable instructions that cause the computer to:

(f) for each condition with a stored identifier that is a join condition, perform step (g);

(g) if the condition is useful in co-locating the sources of the join, identify that condition as CPU effective;

(h) execute the query without applying conditions with stored identifiers that are not identified as CPU effective.

21. A database system for executing database queries, comprising:

one or more nodes;

a plurality of CPUs, each of the one or more nodes providing access to one or more CPUs;

a plurality of virtual processes, each of the one or more CPUs providing access to one or more virtual processes;

each virtual process configured to manage data, including rows organized in tables, stored in one of a plurality of data-storage facilities; and

an optimizer configured to:

(a) identify a first set of conditions corresponding to a selected step for executing a query;

(b) identify a second set of conditions corresponding to one or more steps for executing the query that feed the selected step;

(c) for each condition in the first set, check whether the condition is mathematically redundant, including redundant without being equivalent, with

the other conditions in the union of

the conditions corresponding to the selected step and

the conditions in the second set,

(d) include each condition in the first set that is redundant as checked in step (c) in a third set; and

(e) if there is only one condition in the third set after performing step (d), store an identifier of the one condition.

5 22. The database system of claim 21, wherein the optimizer is further configured to perform a cardinality calculation without considering conditions with stored identifiers.

23. The database system of claim 21, wherein the optimizer is further configured to:

(f) if there are a plurality of conditions in the third set after performing step (d), perform steps (g) and (h);

10 (g) sort the conditions in the third set to identify an initial condition; and

(h) store an identifier of the initial condition.

24. The database system of claim 21, wherein step (e) includes ceasing to store identifiers of additional conditions and the union in step (c) does not include conditions with stored identifiers, and wherein the optimizer is further configured to:

15 (f) if there are no conditions in the third set after performing step (d), cease to store identifiers of additional conditions;

(g) if there are a plurality of conditions in the third set after performing step (d), perform steps (h) and (i);

(h) sort the conditions in the third set to identify an initial condition;

20 (i) store an identifier of the initial condition;

(j) modify the first set of conditions to include only the conditions in the third set that are not the initial condition; and

(k) start again at step (c).

25. The database system of claim 24, wherein the optimizer is further configured to perform a cardinality calculation without considering conditions with stored identifiers.

26. The database system of claim 21 where the one or more steps for executing the query that feed the selected step include all the steps that feed the selected step.

5 27. The database system of claim 21 wherein the optimizer is further configured to execute the query without applying conditions with stored identifiers.

28. The database system of claim 21 wherein the optimizer is further configured to:

(f) for each condition with a stored identifier that is a single table condition, perform steps (g)-(i);

10 (g) if the condition is an equality with an index of the single table, identify that condition as CPU effective;

(h) if the condition is a range comparison of a value ordered index of the single table, identify that condition as CPU effective;

15 (i) if the condition reduces the size of a temporary result in the execution of the query, identify that condition as CPU effective; and

(j) execute the query without applying conditions with stored identifiers that are not identified as CPU effective.

29. The database system of claim 28 where identifying a condition as CPU effective comprises not identifying the condition as CPU redundant.

20 30. The database system of claim 21 wherein the optimizer is further configured to:

(f) for each condition with a stored identifier that is a join condition, perform step (g);

(g) if the condition is useful in co-locating the sources of the join, identify that condition as CPU effective;

25 (h) execute the query without applying conditions with stored identifiers that are not identified as CPU effective.